Bounding Stochastic Dependence, Complete Mixability of Matrices, and Multidimensional Bottleneck Assignment Problems

Utz-Uwe Haus

Institute for Operations Research, Dept. of Mathematics, ETH Zurich, Switzerland, uhaus@ifor.math.ethz.ch

arXiv:1407.6475v1 [math.OC] 24 Jul 2014

Abstract

We call a matrix completely mixable if the entries in its columns can be permuted so that all row sums are equal. If it is not completely mixable, we want to determine the smallest maximal and largest minimal row sum attainable. These values provide a discrete approximation of of minimum variance problems for discrete distributions, a problem motivated by the question how to estimate the α-quantile of an aggregate random variable with unknown dependence structure given the marginals of the constituent random variables. We relate this problem to the multidimensional bottleneck assignment problem and show that there exists a polynomial 2-approximation algorithm if the matrix has only 3 columns. In general, deciding complete mixability is $\mathcal{NP}$-complete. In particular the swapping algorithm of Puccetti et al. [1] is not an exact method unless $\mathcal{NP} \subseteq \mathcal{ZPP}$. For a fixed number of columns it remains $\mathcal{NP}$-complete, but there exists a PTAS. The problem can be solved in pseudopolynomial time for a fixed number of rows, and even in polynomial time if all columns furthermore contain entries from the same multiset.

Keywords: risk aggregation, VaR-bounds, model uncertainty, positive dependence, bottleneck assignment
2000 MSC: 91B30, 05A05, 91B30, 62P05

1. Introduction

The problem we are considering is the following: Given a matrix $A \in \mathbf{R}^{m\times d}$, we are interested in the best way of permuting entries in each column (independently) so that the maximal row sum is minimized, or so that the minimal row sum is maximized. Given d permutations $\Pi = (\pi_1, \ldots, \pi_d) \in \mathfrak{S}(m)^d$ we denote by A^Π the matrix obtained from A by permuting column j by π_j, i.e. $A^\Pi_{i,j} = A_{\pi_j^{-1}(i),j}$. The optimization problem is then

$$\gamma(A) := \min_{\Pi\in\mathfrak{S}(m)^d} \max_{1\leqslant i\leqslant m} \left\{ \sum_{j=1}^{d} A^\Pi_{i,j} \right\} \tag{1}$$

and

$$\beta(A) := \max_{\Pi\in\mathfrak{S}(m)^d} \min_{1\leqslant i\leqslant m} \left\{ \sum_{j=1}^{d} A^\Pi_{i,j} \right\}. \tag{2}$$

We note that aggregation operations other than $+$ are conceivable (e.g., $\min, \max, \times$), but will not be treated here.

This problem is motivated by an application in quantitative finance, but in fact arises whenever one needs to estimate the influence of stochastic dependence on a statistical problem: Consider an aggregate random variable L of the form $L = \sum_{i=1}^{d} L_i$, where the random variables L_i are possibly not independent. Denote by $F_L(x) = P(L \leqslant x)$ the distribution function of L. We are interested in computing the α-quantile (Value-at-Risk, VaR_α) $F_L^{-1}(\alpha) = \inf\{x \in \mathbf{R} : F_L(x) \geqslant \alpha\}$, for $\alpha \in (0,1)$. Often we have no data on the joint distribution L, but only on the marginal distributions F_j of the constituent random variables L_j, and we also lack information on the dependence structure between them.

In the following we will assume that the marginal distributions are discrete, or have been approximated from below and from above as described in [1]: For F_i the generalized inverse is $F_j^{-1}(\alpha) = \sup\{x \in \mathbf{R} : F_j(x) \leqslant \alpha\}$. Consider a discretization in $N+1$ points. Compute the values $q_r^j = F_j^{-1}(r/N)$ for $r \in \{0, 1, \ldots, N\}$. Denoting by $1_{[a,b)}$ the characteristic function on the interval $[a,b)$,

$$\underline{F_j}(x) = \frac{1}{N}\sum_{r=0}^{N-1} 1_{[q_r^j,+\infty)}(x) \text{ and } \overline{F_j}(x) = \frac{1}{N}\sum_{r=1}^{N} 1_{[q_r^j,+\infty)}(x),$$

provide discrete approximations of F_j with $\underline{F_j} \geqslant F_j \geqslant \overline{F_j}$.

Dependence among the individual F_j will manifest itself in the way the values $q_r^j = F_j^{-1}(r/N)$ are appearing in the matrix $A = \begin{pmatrix} q_0^1 & \cdots & q_0^d \\ \vdots & & \vdots \\ q_N^1 & \cdots & q_N^d \end{pmatrix}$. In particular, the row sums may vary significantly: Consider $d = 2$ and the uniform discrete distribution on $\{0, \ldots, N\}$. If L_1 and L_2 are comonotonic (i.e. there is perfect positive dependence among the random variables), then $(q_0^1, \ldots, q_N^1) = (q_0^2, \ldots, q_N^2)$ with row sums $\{0, 2, \ldots, 2N\}$. If, on the other hand, F_1 and F_2 are countermonotonic (perfect negative dependence among the random variables), then $(q_0^1, \ldots, q_N^1) = (q_N^2, \ldots, q_0^2)$, and all row sums are equal to N. If we want to find an upper bound for $F_L^{-1}(\alpha)$ we need to consider

matrices with entries q_r^j for $\frac{r}{N} \geqslant \alpha$, and for lower bounds matrices constructed from q_r^j for $\frac{r}{N} \leqslant \alpha$ and each time minimize the variance of the row sums of A. This intuition is made exact by a representation theorem of Rüschendorf [2, Theorem 2], showing that for discrete distribution functions, and due to the uniform discretization inherent in our definition of $\underline{F_j}$ and $\overline{F_j}$, solving the minimum variance problem amounts to determining $\gamma(A)-\beta(A)$ for the matrix A, since it is enough to minimize over the set of all rearrangements of the F_j. We refer to [1, 3, 4, 5] for recent applications and to [2] and [6, 7] for more details on the general concept of rearrangements of functions.

Example 1 ([8]). Under the Basel II and III regulatory framework for banking supervision, large international banks are allowed to come up with internal models for the calculation of risk capital. For operational risk the so-called Loss Distribution Approach gives them full freedom concerning the stochastic modeling assumptions used. The resulting risk capital must correspond to a 99.9%-quantile of the aggregated loss data over a year. This corresponds to computing the Value-at-Risk $\mathrm{VaR}_{0.999}(L)$ at $\alpha = 0.999$ for an aggregate loss random variable $L = \sum_{i=1}^{d} L_i$, but makes no requirements on the interdependence between the individual loss random variables L_i corresponding to the indivdual business lines: Assumptions made in the calculation must only be plausible and well founded. Estimating the upper bound and lower bound of the VaR over all possible dependence structures is hence relevant both from the regulator's point of view, as well as from the bank's point of view, to estimate worst case hidden risks in the models presented under the Loss Distribution Approach.

Besides computing (or approximating) $\gamma(A)$ and $\beta(A)$, one is also interested in deciding whether for a given matrix $\gamma(A) = \beta(A)$. We will call such a matrix *completely mixable*, in analogy with the definition of this concept by Wang and Wang [4] for distribution functions.

In this paper we show that deciding complete mixability is a strongly $\mathcal{NP}$-complete problem, even for a fixed number of columns, but can be solved using dynamic programming in pseudopolynomial time for a fixed number of rows. We show that the algorithm proposed by Puccetti et al. in [1] to compute $\gamma(A)$ and $\beta(A)$ is not an exact method unless $\mathcal{NP} \subseteq \mathcal{ZPP}$, despite its impressive computational success [8]. Finally, for matrices in fixed (column) dimension we present a polynomial-time approximation scheme.

2. Complexity

It is known that for two columns the complete mixability problem is solvable explicitly (see the references in [9]). This is also apparent by recognizing that the computation of $\gamma(A)$ can be understood as solving a multidimensional bottleneck assignment problem. The multidimensional bottleneck assignment problem asks for the computation of

$$\min_{\pi_1,\ldots,\pi_d} \max_{1\leqslant i\leqslant m} c_{\pi_1(i),\ldots,\pi_d(i)}$$

for a $\underbrace{m \times \cdots \times m}_{d}$ cost table C. Defining $c_{i_1,\ldots,i_d} = A_{i_1,1} + \cdots + A_{i_d,d}$ we see that $\gamma(A)$ can be computed by solving a multidimensional bottleneck assignment problem. Using Observation 1 below we can similarly compute $\beta(A)$ and thus check complete mixability.

In dimension 2, the bottleneck assignment problem models the following problem: Given a set of workers and a set of tasks, where the time of worker i performing task j is c_{ij}, find a simultaneous assignment of all workers to all tasks such that the maximal time spent by any worker (the bottleneck of the schedule) is minimized. Fulkerson et al. showed that the 2-dimensional bottleneck assignment problem can be transformed into a linear assignment problem [10], and thus is polynomially solvable.

The multi-dimensional bottleneck assignment problem of assigning (equal-sized) crews of workers to (equal-sized) groups of tasks is much harder. Even restricted versions of the 3-dimensional version do not admit a polynomial time approximation scheme [11].

By adding $\mu = -\min_{1\leqslant i\leqslant m, 1\leqslant j\leqslant d} A_{ij}$ to each entry of A we can always shift the matrix to make the smallest entry equal to zero, changing all row sums by $+\mu\cdot d$. For convenience we will hence restrict our attention to integral, nonnegative matrices. Assuming integrality is not a major restriction, since rational matrices can without loss of generality be scaled to become integral, and rational matrices provide a dense subset of the real matrices that could arise in discretizing distribution functions.

First note that β and γ are related as follows:

Observation 1. *Let $A \in \mathbf{Z}^{m\times d}$, and $l := \max_{1\leqslant i\leqslant m, 1\leqslant j\leqslant d} A_{ij}$ its largest entry. Define A' by $A'_{ij} = l - A_{ij}$. Then $\beta(A) = d\cdot l - \gamma(A')$.*

Hence we only ever need to consider one of the two values. To see that deciding complete mixability of A and computing β or γ are actually polynomially equivalent we only need the following obvious necessary condition that will also prove useful later on.

Observation 2. *Let $A \in \mathbf{Z}^{m\times d}$. A is completely mixable if and only if $\gamma(A) = \beta(A) = \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{d} A_{ij}$.*

It turns out that this is sufficient for showing linear time decidability of complete mixability if the entries of A are restricted to at most two values: Those can be mapped to $\{0,1\}$, and then the algorithm used in the proof below provides a linear time check for complete mixability:

Theorem 1. *Let $A \in \{0,1\}^{m\times d}$. A is completely mixable if and only if $m \mid \sum_{1\leqslant i\leqslant m, 1\leqslant j\leqslant d} A_{ij}$. The permutation achieving the complete mix can be computed in linear time $O(m\cdot d)$.*

Proof. "$\Rightarrow$" Let $s = \sum_{i=1}^{m}\sum_{j=1}^{d} A_{ij}$. If $m \nmid s$ then A cannot be completely mixable.

"$\Leftarrow$" Assume $m \mid \sum_{1\leqslant i\leqslant m, 1\leqslant j\leqslant d} A_{ij}$. We need to permute the columns of A such that exactly $\frac{s}{m} \in \{0,\ldots,d\} = r$ entries in each row have value 1.

This can always be done: Define for $i \in \{1,\ldots,m\}$ the *defect* $\delta(i) = r - \sum_{j=1}^{d} A_{ij}$ and $\phi = \sum_{i=1}^{m} |\delta(i)|$ the total defect.

Clearly, $\phi = 0$ if and only if all row sums of the matrix are equal to r.

Starting with $j = 2$ define $S_j = \{i \in \{1,\dots,m\} \;:\; \delta(i) > 0, A_{ij} = 1\}$ and $D_j = \{i \in \{1,\dots,m\} \;:\; \delta(i) < 0, A_{ij} = 0\}$. If $S_j \neq \varnothing$ and $D_j \neq \varnothing$ let $t_j = \min\{|S_j|, |D_j|\}$ and swap the entries of column $A_{\cdot j}$ indexed by the largest t_j entries of S_j with those indexed by the smallest t_j entries of D_j. Repeat in increasing order, for all $j \leqslant d$.

Clearly, throughout the procedure the defect of rows with positive defect can only decrease, and the defect of rows with negative defect can only increase; the total defect decreases by $2t_j > 0$ for each swap. Assume that the procedure stops in the last column with a matrix that has nonzero total defect ϕ. Then there must be a row i_1 with positive defect δ_{i_1} and a row i_2 with negative defect δ_{i_2}, since $r = s/m$. Consider some column index l such that $A_{i_1 l} = 1$ and $A_{i_2 l} = 0$. Then the index i_1 was in S_l, and i_2 was in D_l (because the absolute defects of the rows can only have decreased in later steps), but they were not swapped, a contradiction. □

Note that when the algorithm declares A 'not completely mixable', it has computed a permutation achieving maximal row sum.

We note in passing that if $A \in \mathbf{Z}^{m\times d_1}$ and $B \in \mathbf{Z}^{m\times d_2}$ are completely mixable, then so is $-A$ and $(AB) \in \mathbf{Z}^{m\times(d_1+d_2)}$. A more interesting composition is the following:

Propositon 1 (glueing of completely mixable matrices)**.** *Let $A \in \mathbf{R}^{m_1\times d_1}$ and $B \in \mathbf{R}^{m_2\times d_2}$ be completely mixable matrices that have been permuted to each have equal row sums. Then the matrix*

$$A \oplus B = (C_{ij})_{\substack{1\leqslant i\leqslant m_1 m_2 \\ 1\leqslant j\leqslant d_1 d_2}}$$

with $C_{m_2(i-1)+k,d_2(j-1)+l} = A_{ij} + B_{kl}$ (i.e., the block matrix constructed by replacing every entry A_{ij} of A by a block $(A_{ij} + B_{kl})_{\substack{1\leqslant k\leqslant m_2 \\ 1\leqslant l\leqslant d_2}}$) is completely mixable.

Proof. Since A and B have identical row sums σ_A and σ_B (respectively), the row sum of C is always $d_2\cdot\sigma_A + d_1\cdot\sigma_B$, showing complete mixability of C. □

In general checking complete mixability is hard:

Theorem 2. *It is strongly $\mathcal{NP}$-complete to decide whether an integral matrix $A \in \mathbf{Z}^{m\times d}$ is completely mixable. It remains strongly $\mathcal{NP}$-complete for fixed d, and at least weakly $\mathcal{NP}$-complete for fixed m.*

Proof. Even for $d = 3$ we are looking at a Numerical 3-dimensional Matching problem, which is strongly $\mathcal{NP}$-complete [12, problem SP16] (the row sum that needs to be tested is given by Observation 2).

For $m = 2$ we can reduce Number Partition to this problem: Let $(n_1,\dots,n_d) \in \mathbf{Z}^d$ be a multiset of integers, and let $s = \sum_i n_i$. Then $A = \begin{pmatrix} n_1 & \dots & n_d \\ 0 & \dots & 0 \end{pmatrix}$ is completely mixable if and only if $(n_1,\dots,n_d)$ can be partitioned into two multisets of equal size $\frac{1}{2}s$. This problem is known to be (weakly) $\mathcal{NP}$-complete [13]. □

We note that, as is the case for many $\mathcal{NP}$-hard problems, there can not be a polynomial time approximation algorithm computing an approximate value $\gamma'(A)$ that achieves an additive error $|\gamma'(A) - \gamma(A)| \leqslant K$ for some constant K: For a given completely mixable matrix $A \in \mathbf{Z}^{m\times d}$ the matrix obtained by appending the column $(K', 2K', \dots, mK')^\top$ with $K' \geqslant \max 2da^*, K$ (where a^* denotes the largest entry of A) has all row sums separated by at least K', so approximating $\gamma(A)$ to within K amounts to deciding complete mixability.

Clearly, when both d and m are fixed the problem is trivial by enumeration. For fixed m and variable d a dynamic programming algorithm similar to the one for Number Partition of Garey and Johnson [12] can be devised to check complete mixability:

Lemma 1. *There is a pseudopolynomial algorithm to decide complete mixability for matrices $A \in \mathbf{Z}^{m\times d}_{\geqslant 0}$ if the number of rows m is fixed.*

Proof. We can enumerate all possible values appearing as row sums as $v_1,\dots,v_N$, with $N \leqslant d\cdot\max_{1\leqslant i\leqslant m, 1\leqslant j\leqslant d} A_{ij}$. Build a dynamic programming table B with Boolean entries $B(i, j, v_1, \dots, v_N)$, where $B(i, j, \dots, r, \dots)$ is True if and only if value r can be constructed as a (partial) row sum in row i with j columns: Iterate over the columns of A succesively and update B using each of the (fixed number of) permutations that can be applied to column j of A. Then A is completely mixable if $B(i, d, \dots, r, \dots)$ is True for all rows i, where r is the target row sum $\frac{1}{m}\sum_{i,j} A_{ij}$. □

The results in [11] for the bottleneck 3-assignment problem with costs defined by distances (**B3AP-per**) yield a 2-approximation for determining $\gamma(A)$ and $\beta(A)$.

Lemma 2. *For $A \in \mathbf{Z}^{m\times 3}_{\geqslant 0}$ there exists a polynomial 2-approximation algorithm for computing $\gamma(A)$.*

Proof. For convenience we will in this proof assume that the matrix A is indexed by (i, j) with $0 \leqslant i \leqslant m-1$ and $0 \leqslant j \leqslant d-1$. We construct an instance of B3AP-per as follows: Let $I = \{0,\dots,3m-1\}$ denote the indices of all elements of A in column-major order, i.e. index $l \in I$ selects element $(\lfloor\frac{l}{3}\rfloor, l \bmod 3)$ of A, and define the sets $R = \{3k+1 \mid k < m\}$, $G = \{3k+2 \mid k < m\}$, and $B = \{3k+3 \mid k < m\}$ such that $I = R \cup G \cup B$. Define $\mathrm{dist}(i, j) = \frac{1}{2}(A_{(\lfloor\frac{i}{3}\rfloor),(i \bmod 3)} + A_{(\lfloor\frac{j}{3}\rfloor),(j \bmod 3)})$. Then dist satisfies the triangle inequality and is symmetric. It does not necessarily satisfy $\mathrm{dist}(i, i) = 0$, so is not a proper metric. Nevertheless, Theorem 1 of [11] holds with the original proof, as only symmetry and triangle inequality are exploited, and $\mathrm{dist}(i, j)$ is only ever evaluated between pairs of different index sets from $\{R, G, B\}$, i.e. $\lfloor\frac{i}{3}\rfloor \neq \lfloor\frac{j}{3}\rfloor$.With our definition of $\mathrm{dist}(i, j)$

$$\begin{aligned} c_{ijk} &= \mathrm{dist}(i, j) + \mathrm{dist}(j, k) + \mathrm{dist}(k, l) \\ &= \tfrac{1}{2}\big((A_{i1} + A_{j2}) + (A_{j2} + A_{k3}) + (A_{k3} + A_{i1})\big) \;, \end{aligned}$$

since costs need only be defined for $i \in R, j \in G, k \in B$. Then determining $\gamma(A)$ is exactly the B3AP-per problem of [11]. □

3. The swapping algorithm

As noted by Puccetti and Rüschendorf [1], it is sometimes easy to check that a matrix can be permuted so as to increase its smallest row sum. We need the following definition:

Definition 1. For $A \in \mathbf{Z}^{m\times d}$ let $A^{[j]}$ denote the matrix obtained from A by dropping its j-th column, i.e. $A^{[j]} = (A_{\cdot 1} \dots A_{\cdot(j-1)} A_{\cdot(j+1)} \dots A_{\cdot d})$.

For $x, y \in \mathbf{Z}^m$ denote by $x \nparallel y$ that x and y are *oppositely ordered*, i.e.there exists a permutation $\pi \in \mathfrak{S}(m)$ such that $x_{\pi(1} \leqslant \dots \leqslant x_{\pi m}$ and $y_{\pi(1)} \geqslant \dots \geqslant y_{\pi(m)}$.

Lemma 3 (Theorem 3.1 of [1])**.** *Let $A \in \mathbf{Z}^{m\times d}$. If there exists a column index j such that $(\sum_l A^{[j]}_{1l}, \dots, \sum_l A^{[j]}_{ml})^\top \nparallel A_{\cdot j}$, then column $A_{\cdot j}$ can be permuted such that opposite ordering is achieved, and the minimal row sum of A does not decrease.*

For completeness we give the following proof.

Proof. Let $(\sum_l A^{[j]}_{1l}, \dots, \sum_l A^{[j]}_{ml})^\top =: x \nparallel y := A_{\cdot j}$. Then there exists a pair of indices i_1, i_2 such that $x_{i_1} \leqslant x_{i_2}$ and $y_{i_1} \leqslant y_{i_2}$. Therefore $x_{i_1} + y_{i_1} \leqslant x_{i_1} + y_{i_2}$ and $x_{i_1} + y_{i_1} \leqslant x_{i_2} + y_{i_1}$. Hence

$$\min\{x_{i_1} + y_{i_1}, x_{i_2} + y_{i_2}\} = x_{i_1} + y_{i_1} \leqslant \min\{x_{i_1} + y_{i_2}, x_{i_2} + y_{i_1}\},$$

and thus swapping $y_{i_1} \leftrightarrow y_{i_2}$ cannot decrease the minimal row sum of A.

We note that if both $x_{i_1} < x_{i_2}$ and $y_{i_1} < y_{i_2}$, and there are no duplicate entries in x and y, then the minimal row sum of A will actually increase by at least 1 if $i = \operatorname{argmin}_{i_1,i_2}\{x_{i_1}+y_{i_1}, x_{i_2}+y_{i_2}\}$ is chosen minimally. □

In [1] this is taken as a rationale to propose the following algorithm:

Algorithm 1 Swapping Algorithm

```
procedure AntisortColumns(A)
    while ∃ j : (Σ_l A^[j]_1l, ..., Σ_l A^[j]_ml)^⊤ ∦ A_·j do
        x ← (Σ_l A^[j]_1l, ..., Σ_l A^[j]_ml)^⊤
        y ← A_·j
        select (i_1, i_2) from {(i_1, i_2) | x_i1 < x_i2 ∧ y_i1 < y_i2}
        swap A_i1 j ↔ A_i2 j
    end while
end procedure
```

It is then stated and confirmed experimentally that running this algorithm on many randomly permuted copies of the matrix A will usually determine very good bounds for $\beta(A)$ and $\gamma(A)$, and is often very fast. In [8] it is admitted that no analytic proof of convergence to the optimum is known, even when randomly permuting the starting matrix, despite the promising practical results. This is to be expected:

Lemma 4. *The swapping algorithm 1 of [1] does not run in expected polynomial time unless $\mathcal{NP} \subseteq \mathcal{ZPP}$.*

Proof. Consider an instance of the complete mixability problem. Apply the swapping algorithm. Assume that the expected number of times that the input matrix has to be randomly permuted before the swapping algorithm correctly decides complete mixability were of polynomial size. Since we have shown in Theorem 2 that the problem is strongly $\mathcal{NP}$-complete this would yield a zero-error probabilistic polynomial time algorithm [14] for all problems in $\mathcal{NP}$. This would imply $\mathcal{NP} \subseteq \mathcal{ZPP}$. □

In fact, the algorithm may terminate with an approximation error of $O(\max_{ij} A_{ij})$ (Lemma 6).

For some matrices, however, Lemma 3 actually guarantees a positive increase of the minimal row sum: As noted at the end of the proof of Lemma 3, swapping entries in a column, say j, to achieve opposite ordering will actually increase the minimal row sum by at least 1, unless there are duplicate entries in j or duplicate row sums in the matrix $A^{[j]}$. This yields

Observation 3. *Let $A \in \mathbf{Z}^{m\times d}$ be a matrix where all columns have m different entries, and for which all $(d-1)$-column submatrices obtained by deleting a single column have the property that for all possible permutation of column entries their m row sums have m distinct values. Then $\gamma(A)$ and $\beta(A)$ can be determined in pseudopolynomial time using the swapping algorithm.*

It is not unlikely that a matrix with entries drawn uniformly at randomly from a large domain with few rows has no duplicate row sums (Lemma 5), but it seems very hard to trace how this probability evolves after a few steps of swapping.

Lemma 5. *Let $A \in \mathbf{Z}^{m\times(d+1)}_{\geqslant 0}$ be a matrix where each column contains m entries drawn uniformly at random from $\{1, \dots, N\}$. Then the probability $p_{\neq}(A)$ for a d-column submatrix of A to have m distinct row sums is*

$$p_{\neq}(A) \geqslant 1 - O(\tfrac{m^2}{N}).$$

Proof. Consider a d-column submatrix M. Each entry of M is a random variable, independently drawn from $\{1, \dots, N\}$. We consider the entries of M drawn from $1, \dots, N$ row by row. Hence the probability of obtaining sum s in one row is $\Pr[M_{i1} + \dots + M_{id} = s] = \frac{p(s,d)}{N^d}$, where $p(s,d)$ is the number of partitions of s into exactly d parts. The probability of not obtaining sum s is $\frac{N^d - p(s,d)}{N^d}$.

Matrix M has m rows; using the binomial distribution formula the probability of obtaining sum s in m one-row trials is thus

$$\begin{aligned}&\Pr[\text{row sum } s \text{ at least twice in } M]\\ &= 1 - \left(\tfrac{N^d - p(s,d)}{N^d}\right)^m - m\tfrac{p(s,d)}{N^d}\left(\tfrac{N^d - p(s,d)}{N^d}\right)^{m-1}.\end{aligned}$$

Therefore the probability for M to have distinct row sums is

$$\begin{aligned}&\Pr[M \text{ has distinct row sums}]\\ &= 1 - \Pr[\exists s : \text{ row sum } s \text{ at least twice in } M]\end{aligned}$$

and since we can have at most m row sums,

$$\begin{aligned}
&\geqslant 1 - m\Pr[\text{most likely duplic. rowsum } s^* \text{ at least twice in } M] \\
&= 1 - m\left(1 - \left(\frac{N^d - p(s^*,d)}{N^d}\right)^m \right. \\
&\qquad \left. - m\frac{p(s^*,d)}{N^d}\left(\frac{N^d - p(s^*,d)}{N^d}\right)^{m-1}\right)
\end{aligned}$$

where, to upper bound the probability of duplicates, we need to lower bound $(N^d - p(s^*,d))$

$$\begin{aligned}
&\geqslant 1 - m\left(1 - \frac{(N^d - N^{d-1})^m}{N^{dm}} \right. \\
&\qquad \left. - mp(s^*,d)\frac{(N^d - N^{d-1})^{m-1}}{N^{dm}}\right) \\
&\geqslant 1 - m\left(1 - \frac{O(N^{dm} + (-1)^{m-1}N^{d(m-1)})}{N^{dm}} \right. \\
&\qquad \left. - m \cdot 1 \cdot \frac{O(N^{d(m-1)} + (-1)^{m-2}N^{d(m-2)})}{N^{dm}}\right) \\
&\geqslant 1 - O(\tfrac{m^2}{N})
\end{aligned}$$

where for the partition of s into d parts we use the trivial lower bound of 1 and the generous upper bound $p(s,d) \leqslant (s-d+1)^{d-1}$ which is obtained as follows: To partition s^* we need to use at least 1 unit in each of the d parts. We now still can distribute $s^* - d$ units into d bins; we can choose freely from $\{0,\ldots,s^*-d\}$ for $d-1$ bins, then the amount for the last bin is determined. □

4. Matrices of consecutive integers

Definition 2. Let $d, N \in \mathbf{Z}_{\geqslant 0}$ and $a = (1,\ldots,N)^\top$. Every matrix A^Π obtained through permutations $\Pi \in \mathfrak{S}(N)^d$ of the columns from $A = (a,\ldots,a) \in \mathbf{Z}^{N\times d}$ will be called (N,d)-*complete consecutive integers matrix*.

We will now show that for such matrices and certain choices of N (given d) the values of β and γ can be computed explicitly, and that these yield bounds for arbitrary values of N. Furthermore, we will demonstrate that the swapping algorithm of [1] (Algorithm 1) on these instances does not have a constant factor approximation guarantee (it is at least $O(N)$).

Theorem 3. *Let $A \in \mathbf{Z}_{\geqslant 0}^{N\times d}$ be a (N,d)-complete consecutive integers matrix and $N = d^k$ for some $0 < k \in \mathbf{Z}$. Then A is completely mixable and*

$$\gamma(A) = \beta(A) = d + \sum_{i=0}^{d-1}\sum_{j=1}^{k} i \cdot d^{j-1} =: a_d(k).$$

Proof. For $k = 1$ the matrix $A = \begin{pmatrix} 1 & 2 & \ldots & d \\ \vdots & \vdots & \iddots & \vdots \\ d-1 & d & & d-2 \\ d & 1 & \ldots & d-1 \end{pmatrix}$ is a permutation that shows that the (d,d)-complete consecutive integers matrix is completely mixable with uniform row sum $\sum_{i=0}^{d} i = d + \sum_{i=0}^{d-1} i = d + \sum_{i=0}^{d-1} i\sum_{j=1}^{1} d^0 = a_d(1)$.

Assume that the statement holds for $k \in \mathbf{Z}_{\geqslant 0}$, i.e. that a (d^k,d)-complete consecutive integers matrix A of size $d^k \times d$ has been reordered into a matrix A' with identical row sums $a_d(k)$. We will use A' to construct a matrix A'' with d^{k+1} rows that is a reordering of the (d^{k+1},d)-complete consecutive integers matrix of size d^{k+1} and has row sums $a_d(k+1)$: We use the glueing operation of Proposition 1 between A' and $B = d^k\begin{pmatrix} 0 & 1 & \ldots & d-1 \\ 1 & 2 & \iddots & 0 \\ \vdots & \iddots & \iddots & \vdots \\ d-1 & 0 & \ldots & d-2 \end{pmatrix}$ (which has constant row sum $d^k\frac{d(d-1)}{2}$), to obtain $A'' = A' \oplus B$, which has row sum $a_d(k) + \sum_{0\leqslant i<d} i\cdot d^k = a_d(k+1)$. □

Corollary 1. *Let $A \in \mathbf{Z}_{\geqslant 0}^{N\times d}$ be a (N,d)-complete consecutive integers matrix. Then*

$$a_d(\lfloor\log_d(N)\rfloor) \leqslant \beta(A) \leqslant \gamma(A) \leqslant a_d(\lceil\log_d(N)\rceil).$$

In particular, by underestimating $\beta(A)$ as $a_d(\lfloor\log_d(N)\rfloor)$ and overestimating $\gamma(A)$ as $a_d(\lceil\log_d(N)\rceil)$ we make an additive error of at most $\sum_{i=0}^{d-1} i\cdot d^{\lceil\log_d(N)\rceil-1}$ (which is roughly $\frac{d^2N}{2}$).

Lemma 6. *Let A be a $(N,3)$-complete consecutive integers matrix where all permutations are the identity. Then the swapping algorithm will terminate after one reordering step with a matrix with row sums in the range of $[N+2,\ldots,2N+1]$. In particular, if $N = 3^k$ and A is hence completely mixable the solution is has additive error $O(N)$.*

Proof. Starting with $A = \begin{pmatrix} 1 & 1 & 1 \\ \vdots & \vdots & \vdots \\ N & N & N \end{pmatrix}$ the swapping algorithm will invert the order of the first column to obtain $A' = \begin{pmatrix} N & 1 & 1 \\ \vdots & \vdots & \vdots \\ 1 & N & N \end{pmatrix}$. This matrix satisfies the rule that each column is sorted anti-monotonously wrt. the sums of the other two columns, so the algorithm stops. The row sums are $N+2, N+3,\ldots,2N,2N+1$.

Since for $N = 3^k$ we know that there exists a reordering of A such that all row sums are $3 + \sum_{i=1}^{k} 3^i$ this shows an approximation error of at least $O(N)$. □

5. Matrices with restricted domain

Matrices of consecutive integer entries are just a special case of matrices where all columns contain the same multiset of entries $M = \{v_1,\ldots,v_m\}$. If the number of different entries in M is fixed, these matrices yield tractable instances for variable d, much like an N-fold system.

Lemma 7. *Let $A \in \mathbf{Z}^{m\times d}$ such that the entries of each column come from the same multiset $M = \{a_1,\ldots,a_m\}$, and assume m is fixed. Then $\gamma(A)$ can be computed in polynomial time.*

Proof. Since the multiset M is fixed, there are only a fixed number of different ways to rearrange a column by permutations. For each of these k arrangements of the set M denote the permutation by π_l, $1 \leqslant l \leqslant k$. Then

$$\begin{array}{rl} & \min \Gamma \\ \begin{pmatrix} v_{\pi_1(1)} \\ \vdots \\ v_{\pi_1(m)} \end{pmatrix} x_{\pi_1} + \cdots + \begin{pmatrix} v_{\pi_k(1)} \\ \vdots \\ v_{\pi_k(m)} \end{pmatrix} x_{\pi_k} & \leqslant \begin{pmatrix} \Gamma \\ \vdots \\ \Gamma \end{pmatrix} \\ \sum_{l=1}^{k} x_{\pi_l} & = d \\ x_{\pi_l} \in \mathbf{Z}_{\geqslant 0} & \text{for } 1 \leqslant l \leqslant k \end{array}$$

is an integer programming problem in fixed dimension k, modeling that we have to choose d rearrangements of the set M (one for each column of A) that can be solved in polynomial time [15]. □

Instead of instances with the same multiset of values in every column we can also consider instances where all matrix entries come from a fixed set of values, generalizing the two-value case of Lemma 1.

Theorem 4. *Let $M = \{v_1, \ldots, v_s\} \subseteq \mathbf{R}$ be a fixed set of values and $A \in M^{m \times d}$. For every fixed number of columns d one can compute $\gamma(A)$ in polynomial time.*

Proof. If M is fixed then for fixed d there are at most s^d possible row vectors $r_1, \ldots, r_{s^d}$ composed of values from M. We define the binary value u_{ij}^k to be 1 if and only if $(r_k)_j = v_i$, i.e. if in row vector k the value v_i appears in the j-th column.

For a given matrix $A \in M^{m \times d}$ we can count the number of occurences of value v_i in column j in polynomial time. Denote these values by o_{ij}.

Introduce binary variables $p_1, \ldots, p_{s^d}$ to indicate whether pattern k occurs in the permuted version of A, and integer variables $q_1, \ldots, q_{s^d}$ counting how often it appears. Then the following integer program in fixed dimension s^d can be used compute $\gamma(A)$:

$$\begin{array}{rll} \min \Gamma & & \\ (\sum_{j=1}^{d} (r_k)_j) p_k & \leqslant \Gamma & \text{for all } k \\ \sum_{k=1}^{s^d} p_k & \leqslant m & \\ p_i & \leqslant q_i & \text{for all } k \\ q_i & \leqslant m p_i & \text{for all } k \\ \sum_{k=1}^{s^d} u_{ij}^k q_k & = o_{ij} & \text{for all } i, j \\ & p_i \in \{0,1\}^{s^d}, q_i \in Z^{s^d} & \end{array}$$

□

Corollary 2. *There exists a polynomial approximation scheme for every fixed d to compute $\gamma(A)$ for $A \in \mathbf{R}^{m \times d} \geqslant 0$ with multiplicative error $(1+\epsilon)$ for every $\epsilon > 0$.*

Proof. Define a grid of width $\epsilon \frac{a^*}{d}$ where a^* is the largest entry of A. Consider the set $M = \{0, \epsilon \frac{a^*}{d}, 2\epsilon \frac{a^*}{d}, \ldots, \lceil \frac{d}{\epsilon} \rceil \epsilon \frac{a^*}{d}\}$ and round the entries of A up to next value in M to obtain an approximating instance $\bar{A}$. Then by Lemma 4 the approximating instance can be solved in polynomial time since M has $\lceil \frac{d}{\epsilon} \rceil + 1$ entries, a number only depending on the fixed d and ϵ. The objective value of the approximate solution is at most $d\epsilon \frac{a^*}{d} \leqslant \epsilon \gamma(A)$ larger than $\gamma(A)$, since $a^* \leqslant \gamma(A)$, yielding a $(1+\epsilon)$-approximation. □

Acknowledgments

The author wants to thank Giovanni Puccetti for bringing the question to his attention, and David Adjiashvili, Robert Weismantel and Sandro Bosio for helpful discussions.
Part of this research was supported by EU-FP7-PEOPLE project 289581 'NPlast'.

References

References

[1] G. Puccetti, L. Rüschendorf, Computation of sharp bounds on the distribution of a function of dependent risks, Journal of Computational and Applied Mathematics 236 (7) (2012) 1833–1840. doi:10.1016/j.cam.2011.10.015.

[2] L. Rüschendorf, Solution of a statistical optimization problem by rearrangement methods, Metrika 30 (1983) 55–61. doi:10.1007/BF02056901.

[3] G. Puccetti, B. Wang, R. Wang, Advances in complete mixability, Journal of Applied Probability 49 (2) (2012) 430–440. doi:10.1239/jap/1339878796.

[4] B. Wang, R. Wang, The complete mixability and convex minimization problems with monotone marginal densities, Journal of Multivariate Analysis 102 (2011) 1344–1360. doi:10.1016/j.jmva.2011.05.002.

[5] G. Puccetti, L. Rüschendorf, Bounds for joint portfolios of dependent risks, Statistics & Risk Modeling with Applications in Finance and Insurance 29 (2) (2012) 107–132. doi:10.1524/strm.2012.1117.

[6] L. Rüschendorf, Ordering of distributions and rearrangement of functions, Annals of Probability 9 (2) (1981) 276–283.

[7] P. W. Day, Rearrangements of measurable functions, PhD thesis, California Institute of Technology, Pasadena, California, USA (April 1970).

[8] P. Embrechts, G. Puccetti, L. Rüschendorf, Model uncertainty and VaR aggregation, Journal of Banking & Finance 37 (8) (2013) 2750 – 2764. doi:10.1016/j.jbankfin.2013.03.014.

[9] L. Rüschendorf, On the multidimensional assignment problem, Zeitschrift für Operations Research, Serie A 47 (1983) 107–113.

[10] D. Fulkerson, I. Glicksberg, O. Gross, A production line assigment problem, RAND Research Memorandum RM-1102, The RAND Corporation (1953).

[11] D. Goossens, S. Polyakovskiy, F. Spieksma, G. Woeginger, The approximability of three-dimensional assignment problems with bottleneck objective, Optimization Letters 4 (2010) 7–16. doi:10.1007/s11590-009-0152-7.

[12] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-completeness, W. H. Freeman and Company, New York, NY, 1979.

[13] M. R. Garey, D. S. Johnson, "strong" NP-completeness results: Motivation, examples, and implications, Journal of the ACM 25 (1978) 499–508. doi:10.1145/322077.322090.

[14] J. Gill, Computational complexity of probabilistic turing machines, SIAM Journal on Computing 6 (4) (1977) 675–695. doi:10.1137/0206049.

[15] H. W. Lenstra, Integer programming with a fixed number of variables, Mathematics of Operations Research 8 (1983) 538–548.